

2

DTIC FILE COPY

AD-A219 115

IN SEARCH OF INSIGHT

Technical Report AIP - 55

Craig A. Kaplan & Herbert A. Simon

Department of Computer Science
Carnegie Mellon University
Pittsburgh, Pa. 15213

The Artificial Intelligence
and Psychology Project

Departments of
Computer Science and Psychology
Carnegie Mellon University

Learning Research and Development Center
University of Pittsburgh

DTIC
ELECTE
MAR 14 1990
S E D

Approved for public release; distribution unlimited.

90 03 12 015

IN SEARCH OF INSIGHT

Technical Report AIP - 55

Craig A. Kaplan & Herbert A. Simon

Department of Computer Science
Carnegie Mellon University
Pittsburgh, Pa. 15213

15 August 1988

DTIC
ELECTE
MAR 14 1990
S E D

This research was supported in part by the Computer Sciences Division, Office of Naval Research and DARPA under Contract Number N00014-86-K-0678 and the Defense Advanced Research Projects Agency under Contract F-33615-84-K-1520. The authors would like to thank the following people for their useful comments and suggestions regarding earlier versions of this paper: John Anderson, Pat Carpenter, Janet Davidson, Kevin Dunbar, David Klahr, Ken Koedinger, Ken Kotovsky, Deepak Kuilkarni, Jay McClelland, Allen Newell, Mark St. John, David Steier, Ross Thompson and Gregg Yost. Reproduction in whole or in part is permitted for purposes of the United States Government. Approved for public release; distribution unlimited.

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION Unclassified			1b. RESTRICTIVE MARKINGS	
2a. SECURITY CLASSIFICATION AUTHORITY			3. DISTRIBUTION / AVAILABILITY OF REPORT Approved for public release; Distribution unlimited	
2b. DECLASSIFICATION / DOWNGRADING SCHEDULE				
4. PERFORMING ORGANIZATION REPORT NUMBER(S) AIP - 55			5. MONITORING ORGANIZATION REPORT NUMBER(S)	
6a. NAME OF PERFORMING ORGANIZATION Carnegie-Mellon University		6b. OFFICE SYMBOL (if applicable)		7a. NAME OF MONITORING ORGANIZATION Computer Sciences Division Office of Naval Research
6c. ADDRESS (City, State, and ZIP Code) Department of Psychology Pittsburgh, Pennsylvania 15213		7b. ADDRESS (City, State, and ZIP Code) 800 N. Quincy Street Arlington, Virginia 22217-5000		
8a. NAME OF FUNDING / SPONSORING ORGANIZATION Same as Monitoring Organization		8b. OFFICE SYMBOL (if applicable)		9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER N00014-86-K-0678
8c. ADDRESS (City, State, and ZIP Code)		10. SOURCE OF FUNDING NUMBERS p4000ub201/7-4-86		
		PROGRAM ELEMENT NO N/A	PROJECT NO N/A	TASK NO. N/A
		WORK UNIT ACCESSION NO N/A		
11. TITLE (Include Security Classification) IN SEARCH OF INSIGHT				
12. PERSONAL AUTHOR(S) Craig A. Kaplan & Herbert A. Simon				
13a. TYPE OF REPORT Technical		13b. TIME COVERED FROM 86Sept15 to 91Sept14		14. DATE OF REPORT (Year, Month, Day) 1988 August 15
				15. PAGE COUNT 55
16. SUPPLEMENTARY NOTATION				
17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)	
FIELD	GROUP	SUB-GROUP		
			Insight, problem solving, heuristic search.	
19. ABSTRACT (Continue on reverse if necessary and identify by block number)				
<p>This paper describes the process of attaining insight in the domain of a particular insight problem -- the Mutilated Checkerboard (MC) problem. Specifically, it shows that the process of attaining insight can be viewed as search, and that performance on insight problems can be predicted by the availability of sources of search constraint. To test these claims we conducted an experiment that varied the salience of features leading to the critical concept of parity in the MC problem. Using chronometric measures, analyses of verbal protocols, and computer simulation techniques, we explored first the reason for the difficulty of the Checkerboard problem, and then four potential sources of search constraint. Results concerning the effects of cue salience manipulations, prior knowledge, hints, and use of heuristics are presented. While subjects used each of these four sources of constraint, noticing properties of the situation that remained invariant during solution attempts (the Notice Invariants heuristic) was a very powerful means for focusing search for a viable problem space. We show that, in conjunction with hints and independently, it played a major part in producing insight into the solution.</p>				
20. DISTRIBUTION / AVAILABILITY OF ABSTRACT <input type="checkbox"/> UNCLASSIFIED/UNLIMITED <input checked="" type="checkbox"/> SAME AS RPT <input type="checkbox"/> DTIC USERS			21. ABSTRACT SECURITY CLASSIFICATION	
22a. NAME OF RESPONSIBLE INDIVIDUAL Dr. Alan L. Meyrowitz			22b. TELEPHONE (Include Area Code) (202) 696-4302	22c. OFFICE SYMBOL N00014

DD FORM 1473, 84 MAR

83 APR edition may be used until exhausted.

All other editions are obsolete.

SECURITY CLASSIFICATION OF THIS PAGE

Unclassified

ABSTRACT

This paper describes the process of attaining insight in the domain of a particular insight problem -- the Mutilated Checkerboard (MC) problem. Specifically, it shows that the process of attaining insight can be viewed as search, and that performance on insight problems can be predicted by the availability of sources of search constraint. To test these claims we conducted an experiment that varied the salience of features leading to the critical concept of parity in the MC problem. Using chronometric measures, analyses of verbal protocols, and computer simulation techniques, we explored first the reason for the difficulty of the Checkerboard problem, and then four potential sources of search constraint. Results concerning the effects of cue salience manipulations, prior knowledge, hints, and use of heuristics are presented. While subjects used each of these four sources of constraint, noticing properties of the situation that remained invariant during solution attempts (the Notice Invariants heuristic) was a very powerful means for focusing search for a viable problem space. We show that, in conjunction with hints and independently, it played a major part in producing insight into the solution.

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	



Most of us have experienced insight while trying to solve a problem. Subjectively, we might speak of an AHA! experience in which we suddenly felt that we knew the answer to a problem, even if we could not specify its details. While opinions vary as to a precise definition of insight (Weisberg & Alba 1981a, 1981b, 1982, Dominowski 1981, Ellen 1982, Ohlsson 1984a), many researchers have mentioned the subjective AHA! feeling as a critical component (Posner 1971, Kohler 1956, 1969, Duncker 1945, Worthy 1975). Moreover, the AHA! experience figures prominently in anecdotal accounts of insightful discoveries (Haefele 1962, Ghiselin 1952, Hadamard 1949). For the purposes of this paper, we will follow the convention of using insight to refer to a subjective AHA! experience during problem solving.

There are many potential causes of surprise in problem solving. However, we will focus on insight resulting from a change in representation. Such changes seem to typify many of the problems used in past experimental studies. For example, subjects in Duncker's functional fixedness experiments must re-represent the function of a critical object before they are able to reach a solution. There can be little doubt that Duncker considered this "restructuring" to be the essential source of difficulty (as well as of the Aha! experience) when he wrote (Duncker 1945, pg. 29):

It has often been pointed out that such restructurations play an important role in thinking, in problem solving. The decisive points in thought processes, the moments of sudden comprehension, of the "Aha!," of the new, are always at the same time moments in which such a sudden restructuring of the thought material takes place, in which something "tips over."

A second example of insight co-occurring with representational change can be found in Kohler's experiments with chimpanzees. The famous account of the way in which the ape Sultan insightfully joins two sticks to form a pole long enough to reach some bananas outside his cage surely involves representational change (Kohler 1956). Before his insight, Sultan knew he could use a single stick as a pole, and in fact Kohler tells us that Sultan tries to reach the fruit this way but fails. After his insight, Sultan has clearly acquired the concept of joining two sticks to make one, and Kohler tells us that Sultan was able to use this method on subsequent trials with little delay. This change in Sultan's representation of the potential uses for a stick is quite sudden, and seems to be accompanied by the chimpanzee equivalent of an AHA! experience (Kohler 1956, pg 127.):

Sultan first of all squats indifferently on the box, which has been left standing a little back from the railings; then he gets up, picks up the two sticks, sits down again on the box and plays carelessly with them. While doing this, it happens that he finds himself holding one rod in either hand in such a way that they lie in a straight line; he pushes the thinner one a little way into the opening of the thicker, jumps up and is already on the run towards the railings, to which he has up till now half turned his back, and begins to draw a banana towards him with the double stick.

Similar arguments can be made in the case of the nine dots problem (Lung & Dominowski 1985, Weisberg & Alba 1981, Burnham & Davis 1969), the match stick problems (Katona 1940), and the two string problem (Maier 1931). Most of the remainder of this paper will address the issue of representational change in yet another insight problem, the Mutilated Checkerboard problem.

An important motivation for this research is the belief that insight can be best understood by understanding the processes that underlie it. It is in this spirit that we begin our investigation with a process oriented metaphor.

Diamonds In the Dark

Imagine that you are searching for a diamond in a huge dark room. The room may contain a light switch, but neither the diamond nor the light switch are located where you expected them to be. What do you do?

We believe that subjects asked to solve insight problems are faced with an analogous task. In insight problems, the "light switch" is a particular way of looking at the problem, a critical representation, that makes the nature of the solution apparent.¹ Most insight problems are difficult primarily because the solver is "in the dark" with regard to the critical representation. In contrast to routine problems where prior experience is usually quite helpful in achieving a rapid solution, insight problems often have the property that past experience misleads rather than helps. Despite these obstacles, most people are somehow able to solve insight problems, with varying degrees of efficiency.

The aim of this paper is to describe the *process* of attaining insight in the domain of a particular problem. Specifically, we claim that the process of attaining insight can be viewed as search, and that performance on insight problems can be predicted by the availability of sources of search constraint.

To develop some intuitions about our claim, consider some of the possible actions available to our seeker of diamonds in the dark room. One approach would be to explore the room randomly, hoping to bump into either the light switch or the diamond. Analogously, chance seems to have played at least some role in a number of great scientific insights (e.g. the discoveries of X-rays, of the vulcanization of rubber, and of penicillin). However both diamonds and insights would be much rarer than they already are if their discovery depended upon chance alone.

A better strategy might be to restrict search for the diamond to the area of the room that seems most promising. Alternatively one might search for the light switch, rather than the diamond, reasoning that the light switch ought to be easier to find, and that light would make the location of the diamond apparent. Both approaches acknowledge the huge size of the room but try to maximize the chances of finding the diamond by somehow constraining or guiding search. Similarly, we will argue that subjects solve insight problems by constraining their search. Understanding insight, we claim, amounts largely to understanding *the ways in which* subjects constrain their search.

Heuristic Search -- Adding Rigor to the Metaphor

Newell & Simon's (1972) conception of problem solving as heuristic search through a problem space, combined with Simon & Lea's (1974) notion of a dual problem space for instances and hypotheses captures many of our intuitions in a rigorous theory. The dark room maps nicely onto the search space of the task environment -- that is a space of all possible actions that might be taken with regard to a particular problem.

Of course, the actions that a problem solver actually takes depends upon his/her representation of the problem. Each representation corresponds to a problem space, with changes in representation

¹ Note: Depending on the type of problem, a fair amount of problem solving *making use of* the critical representation may still be necessary to complete the problem solution (e.g. the nine dots problem, Weisberg & Alba 1981).

corresponding to changes in problem space. Problem solvers apply operators associated with a given problem space in an attempt to transform their current (physical or mental) state into a state that satisfies their goal.


Within a given problem space, the trick lies in searching for the right operator to apply next. Similarly, if no operators seem to be working within a given problem space, one must search for a new problem space to explore. Both search within a problem space (for the next operator) and search in the meta-level space of possible problem spaces is often enormously difficult unless constraints for this search can be found.

The crux of Newell & Simon's theory of problem solving is the concept of guided search. Humans seem unwilling to search randomly. For example, we know of no one who would try to solve the problem of misplaced car keys by searching in a *completely* random manner. Using heuristics is one way people constrain their search. These heuristics can range from general ("When something is lost, think of the last time you had it") to specific ("I often leave my car keys in the ignition when I go shopping because I need both hands to carry groceries"). Cues in the environment (e.g. a buzzer that sounds when you leave the car keys in the ignition) can also serve to constrain search. The point is, that a wide range of problem solving behavior can be understood by examining the heuristics and other sources of constraint on search. We hope to understand the process of insight by similarly identifying the heuristics and other sources of search constraint that lead to successful performance on a particular insight problem, The Mutilated Checkerboard problem.

The Mutilated Checkerboard (MC) problem

The MC problem has somewhat of a reputation, both as an insight/puzzle problem (Anderson 1985, Wickelgren 1974) and as a challenge to problem solving programs in the AI community (Korí 1980, Newell 1965, McCarthy 1964). Its difficulty has been assumed to stem from the fact that the initial representation that problem solvers are likely to form fails to solve the problem. Subjects need to change their representation in a non-obvious way.

The classic MC problem (see Figure 1) consists of a standard 8X8 checkerboard whose diagonally opposite corners have been removed. Subjects are told to imagine placing dominos on the board so that one domino covers two horizontally or vertically (but not diagonally) adjacent squares. The problem is either to show how 31 dominos would cover the 62 remaining squares, or to prove logically that a complete covering is impossible. (If you have never seen this problem before, you might want to try it now, before reading the solution.)



The Classic Mutilated Checkerboard (MC) problem
FIGURE 1

Since each of the 31 dominos covers two squares, a covering initially seems possible. To see why a complete covering is actually **impossible**, observe that a domino must always cover a black and a white square. Note that removing two squares of the same color (the diagonal corners) from the 8x8 board has left an imbalance between the number of black and white squares that remain. After covering 30 black-white pairs with 30 dominos, the problem solver is always left in the impossible situation of having to cover two same-colored squares with the single remaining domino.

Pilot Work on problem Difficulty

To solve the MC problem insightfully, subjects must switch from an initial representation that considers only the numbers of squares and dominos, to a new representation that takes the parity of the squares into account as well. One way of representing parity is to partition the squares into two equivalence classes: black squares and white squares.² Switching to such a representation allows subjects to reason about the numbers of squares of each type, and to make the crucial inferences needed to solve the problem (Korf 1980). Our initial pilot work explored the assumption that this switch in representation corresponds to the problem's source of difficulty.

We first tried to get a sense of the magnitude of the problem's difficulty by estimating the size of the search space. The initial obvious way to prove the problem impossible is by trying coverings exhaustively. We constructed a computer program incorporating some simple covering heuristics (e.g.

²Note, there is nothing special about black and white except that color is a readily available classification scheme. Other systems, such as labeling squares even and odd, would work equally well.

look ahead one move and try to constrain maximally the number of possible next placements). The program required 758,148 domino placements in order to prove the problem impossible by exhaustion. The size of this initial problem space clearly makes exhaustive search for a covering (or proof of impossibility by exhaustion) an unacceptable strategy for human subjects.

Of course, a clever subject might discover more powerful heuristics that would reduce the search space further, and might be patient enough to make the entire search, but such cleverness and persistence are rare. Allen Newell (personal communication) reports that he did succeed in solving the problem by exhaustion, using search heuristics that reduced the search tree to a few thousand possibilities. Even with this reduction, few subjects would persevere until they had searched the space completely. To solve the problem, subjects must explore other problem spaces.

A somewhat different perspective on the problem's difficulty follows from a pilot experiment involving the MC conducted by Deepak Kulkarni (Kulkarni, personal communication). In this experiment, a graduate student in Chemical Engineering spent 18 hours and filled 61 pages of a lab notebook with notes, yet still did not solve the problem! While the notebook contains numerous drawings of boards and potential domino placements, the boards were never drawn with alternating squares shaded differently. Since the graduate student was given only a written description of the problem, and not an actual checkerboard, presumably the color of the squares was not available to him unless he shaded the squares himself. (The significance of this seemingly minor detail may become apparent when we discuss our own BREAD & BUTTER experiment below.)

The graduate student, although quite persistent, was eventually forced to try other methods besides exhaustive coverings. Many of his boards were labeled with x and y coordinates, and many pages were devoted to mathematical analyses of various sorts (e.g. equations involving the number of squares as a parameter, degrees of freedom tables, etc.) Some of his later attempts to prove the problem impossible, such as the "anti-puzzle" approach,³ indicate that he had abandoned the space of exhaustive coverings and had searched in a meta-level space of possible new approaches to the problem.

In other pilot research with the MC problem, none of our subjects was able to solve the problem within an hour without being given one or more hints. Again, we found strong evidence that subjects switched from searching in the initial covering problem space, to a meta-level space of potential new approaches to the problem. The pilot data also suggested that subjects had little difficulty in generating a rough proof of the problem's impossibility once they noticed that the parity of the squares was important. Most of their time seemed to be spent either fruitlessly trying various coverings or searching for new approaches to the problem. When subjects finally paid attention to the parity of the squares, many experienced a sudden insight leading to the problem's solution.

³His idea was: "The given puzzle is equivalent to the following anti-puzzle: Given a 1×1 chessboard and dominos, arrange the dominos so that two opposite squares are covered and all others are empty. Soln: [There is] no possible way to cover the 2 opposite corner squares on any given chessboard without covering other squares. [Therefore the original checkerboard problem must also be impossible]." To show why this proof fails, one need only remove one square of each color from the edges of the board. His argument still holds, only now the problem *is* possible!

Theory and Experiment

We have argued that the MC problem is hard because solving it requires discovering a representation (problem space) within which the solution can be found easily. The difficulty lies in discovering the problem space, and not in solving the problem once an appropriate space has been found.

We have also demonstrated that the initial problem space provided by the instructions for the MC problem, the space of all possible coverings, is too large to permit humans to demonstrate impossibility by searching for all legal coverings. In fact, none of our subjects solved the problem by search in this space (although most of our subjects persisted in trying to do so for some time).

The solver must therefore undertake a search in another space: the meta-level space of possible problem spaces for the MC problem. The decision to search for a new representation is motivated by the "Try a Switch" meta-heuristic: If at first you don't succeed, search for a different problem space. But the space of possible problem spaces is exceedingly ill-defined, and probably infinite. Hence, a successful search in this space would require the subject to have or obtain strong constraints that guide search and make it highly selective.

Four Sources of Search Constraint

The remainder of this paper constitutes an attempt to explain the process of changing representations in the MC problem using the concept of search constraint as the unifying framework. Based upon our pilot research and other research along these lines (e.g. see Newell, Shaw & Simon 1962), we have identified four major sources of search constraint that seem relevant to the MC problem, and perhaps insight problems in general:

First, features of the problem itself provide cues which can be used to constrain search. Manipulating the salience of these cues might be one way to affect the solver's ability to attain insight. For example, Janet Davidson found that highlighting relevant information improved childrens' performance on a group of insight problems (Davidson 1986). Similarly, if the critical feature of parity in the MC problem could somehow be made more salient, we would expect subjects to switch representations, and discover the reason for the problem's impossibility more easily

A second source of constraints, hints from the experimenter, tell the subject what features of the problem situation are relevant to the problem's solution. Focusing on these features ought to greatly reduce the time spent exploring unproductive problem spaces. Experiments dating at least back to Maier's (1931) famous two string problem testify to the potential effectiveness of even seemingly unsubstantial hints (e.g. the mere "accidental" brushing of a string by the experimenter).

Both hints and cue salience are external sources of search constraint. Because they are to be found in the problem environment, rather than in the problem solver's head, they can be easily manipulated, and the effects can be observed across subjects. Interacting with these external factors however, are two internal sources of search constraint: domain specific prior knowledge and widely applicable heuristic knowledge.

In a sense heuristics are a special type of prior knowledge, namely knowledge about what strategies

have proved useful in the past. However, the distinction can be made between heuristic knowledge that might be applied to a variety of problems, and specific knowledge of particular domains (e.g. see Larkin, Reif, Carbonell, & Gugliotta 1985).

As pointed out earlier, **relevant** domain knowledge can make a problem routine. However, the solver who unwittingly misapplies this knowledge may spend a long time exploring unfruitful paths. For example, from his notebook, it appears that the Chemical Engineering graduate student mentioned earlier spent the majority of his eighteen hours exploring fruitless mathematical approaches. Someone with much less mathematical experience might have exhausted more quickly his stock of possibilities, while someone with slightly different knowledge (e.g. experience with parity problems) may have been able to constrain search in a productive manner.

The use of heuristics promises to be the most interesting source of search constraint since it offers the clearest opportunity for exploring individual differences between problem solvers. Recent research in the domain of Scientific discovery (e.g., see Langley, Simon, Bradshaw, & Zytkow 1987) suggests that features of a problem might be noticed because they remain invariant or recur repeatedly. For example in the MC problem, subjects may notice that both of the deleted squares are of the same color, or that the squares they fail to cover are always of the color opposite to that of the deleted squares.

Note that all of these sources of constraint direct search at the meta-level problem space of possible representations. Cues, hints, prior experience, and heuristics for exploiting regularities in the problem, each guide the solver to a particular representation -- a problem space that can then be explored for a solution to the problem. Some of these sources of constraint also operate at the level of search within a particular representation. However, since insight in the MC problem appears to occur when subjects discover the appropriate representation (as opposed to when they are trying to find the exact formulation of the proof statement) we will focus on constraints upon search operating in the meta-level problem space.

Description of the BREAD & BUTTER Experiment

To test our theory of insight in the domain of the MC problem, we need first to verify that the problem difficulty stems from search for a new representation, and second to examine potential sources of constraints for this search. By manipulating the salience of the parity cue in the MC problem, and by selectively providing hints, we were able to assess the effect of external sources of search constraint on subject's problem solving behavior. A detailed analysis of thinking aloud protocols allowed an analysis of how internal sources of search constraint (i.e. the use of heuristics and domain specific knowledge) affected problems solving.

As shown in Figure 2, subjects in the BREAD & BUTTER experiment received one of four types of checkerboards which varied with respect to the salience of the critical cue, parity. Hints were provided systematically (if needed) to ensure that all subjects solved the problem within an hour. We predicted that subjects in the higher cue salience conditions would take less time and require fewer hints to solve the problem than subjects in the lower cue salience conditions. We further hoped to find evidence of the use of heuristics in subjects' verbal protocols. Our specific predictions follow the description of experimental methodology below.

Method

Materials. Instructions were typed on a standard (8.5 x 11 inch) sheet of paper. All subjects received identical instructions (see Appendix B). The instructions referred to the "board" but did not mention color or that the board could be thought of as a checkerboard.

The boards were all 8x8 matrices of 3/4 inch squares. These squares were either left blank, filled in with black and pink color in the fashion of a checkerboard, filled with the words "black" and "pink" in checkerboard fashion, or filled with the words "bread" and "butter" in checkerboard fashion (see Appendix B). Blue X's made of sticky paper were placed on the boards to indicate the squares removed.

No real dominos were used, but subjects were allowed to write on the boards with either a pen or pencil. A cassette tape recorder with a condenser mike was used to record thinking aloud protocols.

Procedure. Each subject was run individually. The subject was asked a series of questions regarding class level, major, and previous problem-solving experience. Next subjects were given practice thinking aloud on an unrelated task (e.g. mental multiplication of two three digit numbers). Question answering and protocol-giving practice took about ten minutes.

During the remaining 50 minutes (the session was limited to approximately one hour), subjects were presented with the instructions and asked to solve the problem. After the subjects had read the instructions, the experimenter carefully placed the blue X's on the upper left and lower right hand corners of the board while the subject watched. The experimenter then sat behind the subject where he could view the subject's behavior but where the subject could not see him without turning around.

The subjects were allowed to work for approximately fifteen minutes -- uninterrupted except for occasional prompts to "keep talking" if there were periods of silence. After this time, if the subject had not yet determined that the problem was impossible (many subjects seemed convinced that they could find a way to make the dominos cover the remaining squares if only they persisted long enough), the experimenter told the subject that the problem was indeed impossible and that effort should be spent in trying to prove logically why this was so -- the IMPOSSIBLE HINT. This information was provided because it appeared quite possible for subjects to spend an entire hour trying to prove the problem possible -- an activity that was not likely to lead to any change in representation.

Roughly twenty minutes after having read the instructions, if subjects were still trying to prove the problem impossible by trying different covering approaches, subject were told that there was a trick way of looking at the problem that did not involve exhaustive covering -- the INSIGHT HINT. Fifteen minutes after the INSIGHT HINT was given, if the subjects still had not solved the problem, subjects were told that the color (or words) on the squares might help them solve the problem -- the PARITY HINT.

The PARITY hint differed slightly depending upon the condition the subject was in. For the BLANK condition, subjects were told to take their pencil and color in every other square in the fashion of a checkerboard. They were then to look at the resulting pattern and see if that might help them solve the problem. Subjects in the COLOR condition were told to look at the colors of the squares in the checkerboard to see if that might lend them some insight. Subjects in the BLACK & PINK and BREAD & BUTTER conditions were told to pay attention to the words that were inside the squares in order to gain an insight.

Fifteen minutes after the PARITY hint, if subjects still had not solved the problem, they were instructed to count the relative number of the different types of squares (e.g. "count the numbers of blacks versus pinks") -- the COUNT HINT. In the rare event that a subject failed to solve the problem after being told to count, increasingly directive hints were given until the problem was solved. The problem was considered solved, when subjects were able to generate a rough proof of the problem's impossibility.

A Rough Proof consisted of a statement by the subject expressing a way to solve the problem using the insightful representation. The statement needed to indicate both a recognition of the importance of parity (e.g. describing the color of the squares or words on them as being crucial to the solution) as well as some conviction that the problem had been solved. Time to Rough Proof was used as the dependent measure since we discovered that subjects varied considerably in their opinions (and knowledge) as to what constituted a more formal proof.


Dependent Measures. Times to Rough Proof and to the first mention of parity (e.g. the color or words on the squares) were derived from the tape recordings of all 23 subjects. Notations were made about each of the 23 tapes, and 8 were selected (to include two subjects from each group -- one above, and one below the median solution time for that group) for complete transcription and detailed analysis. The coding system used for these 8 protocols is detailed in Appendix A.

Subjects . Twenty-five CMU undergraduates fulfilling a course requirement in an Introductory Psychology course served as subjects. The subjects were naive to the degree that they were honest about not having had any previous experience with problems involving checkerboards and dominos. Two subjects were suspected of having seen the problem before and their data has been excluded from the subsequent analyses.


Subjects were arbitrarily assigned to conditions (before they were seen), with 5 to 7 subjects per condition.

- 25 SUBJECTS (23 Naive, 2 Excluded -- suspect prior knowledge)
- 4 CONDITIONS, 5-7 Subjects Per Condition
- HINTS GIVEN AT REGULAR INTERVALS

THE FOUR CONDITIONS:




BLANK




COLOR

← PINK
← BLACK



BLACK & PINK



BREAD & BUTTER

(Note: Boards not drawn to actual size)

PREDICTION:

BLANK > COLOR > BLACK & PINK > BREAD & BUTTER
Most Difficult.....Easiest

Experiment 2 at a Glance

FIGURE 2

Predictions:

First, we predicted that parity⁴ would prove to be a critical cue. Once subjects questioned the possible importance of the parity of the squares, we predicted a change in representation followed rapidly by the AHA! experience. Similarly we predicted that most of the problem's difficulty (as reflected by solution times) would stem from ineffective attempts to cover the board or from search for a better representation of the problem.

Because the BLANK condition offered no visually available pattern for subjects to use as the basis for construction of an elegant representation, we predicted that this condition would be most difficult for subjects. The COLOR condition was predicted to be second in difficulty because actual colors are so much a part of one's normal conception of a checkerboard that they tend to be overlooked. On the other hand, words in place of colors should seem unusual and attract attention, so we predicted that the conditions with words would be easier than the other two. In addition, the words "Bread" and "Butter" seem to "go together", and we thought this particular choice of words might emphasize the concept of a pair. Since the realization that a domino must cover a *pair* of differently labeled squares is crucial to the problem's solution, we predicted that the BREAD & BUTTER condition might be easier than the BLACK & PINK condition.

We predicted that hints would be effective in helping subjects constrain their search, and that hints mentioning the parity pattern would be especially effective given that parity is central to the representation that must be acquired to solve the problem insightfully. We also hoped to find evidence of search (and of the heuristics that guide it) in the verbal protocols.

Results

We have claimed first that the difficulty of the MC problem stems from search, and second that performance on this problem can be predicted by the sources of constraint for this search. Specifically, we predicted that subjects could solve the problem more quickly in the more parity-salient conditions, that hints (especially the PARITY hint) could help subjects, and that heuristics (e.g. the Notice Invariants heuristic) might prove not only to be a major source of search constraint, but also might constitute a source of individual differences.

The evidence marshaled to test these claims is of two types. First there are a number of solution time results as well as simple statistical counts that involve data from all twenty three subjects. Second, detailed protocol analyses have been performed on a more manageable subset of eight protocols. These eight protocols, consisting of one fast and one slow (based on median splits of solution time) subject selected arbitrarily from each of the four conditions, also serve as the basis for our claims regarding individual differences.

⁴In the context of any discussion of the BREAD & BUTTER experiment, we use the term parity to refer to the actual features of the squares that allow them to be divided into two equivalences classes. These features vary of course depending upon the experimental condition. In the COLOR condition, the features would be actual colors of the squares, while in the BLACK & PINK and BREAD & BUTTER conditions "parity" refers to the words labeling the squares.

Search as the Source of Problem Difficulty

Analyses of pilot subjects solving the MC problem suggests that their problem solving behavior can be divided into two segments: searching for an appropriate representation, and reasoning within that representation. In this section we argue that the difficulty stems from search and not from difficulty in making necessary logical inferences once the concept of parity has been represented.

Describing Search

Subjects can be viewed as searching at two levels. When they have a particular representation that they believe will allow them to solve the problem, subjects search within the problem space corresponding to that representation. For example, the task instructions suggest that a simple covering of the board might be possible, leading subjects to search initially within a COVERINGS problem space.

As each covering attempt fails, subjects are forced to search in the meta-level space of potential representations to find their next approach. Once found, this new approach constitutes a new problem space which can then be searched. Table 1 lists a sample of the problem spaces actually used by subjects in the BREAD & BUTTER experiment.

The first five approaches listed, corresponding to small shifts in representation around the theme of covering, are typical of the problem spaces explored early in problem solving. The remaining approaches correspond to the more radical changes in representation that typify later problem solving efforts.

TABLE 1

A Dozen of the Approaches (Problem Spaces) Used by Ss

Try placing all dominos horizontally.
 Try placing all dominos vertically.
 Try placing dominos in a spiral pattern.
 Try placing dominos in a zig-zag pattern.
 Try decomposing board into smaller areas, and cover each area.
 Consider whether a path between the Mutilated squares contains an even or odd number of squares.
 Consider if legal moves in the game of checkers are related to the problem.
 Try to draw an analogy between the Checkerboard problem and the 8-puzzle problem.
 Try solving the problem if a different pair of squares were mutilated.
 Try representing the squares as dots, and dominos as lines connecting dots.
 Try rotating the board to see if that changes one's perspective.
 Consider how color might help solve the problem.

Later, we shall focus on search in the meta-level space of possible representations when we discuss heuristics as a source of search constraint. At this point, we wish only to distinguish search at the two levels, and more importantly to distinguish both types of search from the rapid reasoning behavior that immediately precedes subjects' solutions.

Simulating Rapid Reasoning

Figure 3 presents excerpts from three typical subjects in the BREAD & BUTTER experiment which capture their behavior just before, during, and just after the AHA! experience. Each episode lasts less than a minute and a half, with the actual insight being much more rapid. Each contains sufficient information to constitute a Rough Proof of the problem's impossibility. While there are individual differences in the routes taken to the solution, the subjects all seem to draw a series of rapid inferences directly following their insight. Having represented parity, the generation of a Rough Proof seems almost trivial.

SUBJECT 1 (a BREAD & BUTTER subject):

EXCERPT LASTS: 70 Secs.

- 1: Just by trial and error I can only find 31 places ... I dunno, maybe someone else would have counted the spaces and just said that you could fit 31, but if you try it out on the paper, you can only fit 30.
(pause & distracted chattering)
- E: Keep trying.
- 1: Maybe it has to do with the words on the page? I haven't tried anything with that.
(pause)
Maybe that's it. Ok, dominos, umm, the dominos can only fit ... alright, the dominos can fit over two squares, and no matter which way you put it because it cannot go diagonally, it has to fit over a butter and a bread. And because you crossed out two breads, it has to leave two butters left over so it doesn't ... only 30, it won't fit. Is that the answer?

SUBJECT 2 (a COLOR subject):

EXCERPT LASTS: 48 Secs.

- 2: There's an even number of squares, so it's possible depending on the placement... so it has to be the placement. (pause)
- 2: How about a different placement? We could try that.
Well, if we place the Xs in different corners, then it'd be really simple ... other than opposite ummm How about a black and a pink Oh, we always have to cover a black and a pink square... at the same time time Uh, there's no way to avoid that ... ummm.
- Oh!, There's two black squares covered up and ... since you always have to cover up a black and a pink square, there's no way you can do it.

SUBJECT 7 (a COLOR subject requiring a hint): EXCERPT LASTS: 36 Secs.

- E: What about the color? Can you use color to help you out?
- 7: There's two pinks next to each other Oh God!! You're taking two black out? And you would need to take a black and a white out ... a black and a pink out. (pause)
- 7: So you're leaving ... OH!!! Jeez! So you're leaving it's short -- how many, you're leaving uhhhh there's more pinks than black, and in order to complete it you'd have to connect two pinks but you can't because they are diagonally ... is that getting close? ... since they are diagonally connected ... and so you're always gonna end up with two extra pinks ... because their mates were taken out.
-

The AHA! Experience: 3 Protocol Excerpts

FIGURE 3

One way to specify rigorously the work required to generate a Rough Proof once parity has been hinted at, is to simulate the actual switch of representation and subsequent reasoning processes in a computer program. To this end, we built a computer simulation, SWITCH (Kaplan 1988), using the Soar architecture (Laird, Newell, & Rosenbloom 1987).

What the Simulation Starts With

SWITCH starts with essentially the same information as a subject who has already done a significant amount of unsuccessful problem solving and has just been given a hint to pay attention to the color of squares. In addition to modeling the behavior of the subject, however, SWITCH has the task of modeling the environment in which the subject acts. These two sources of knowledge -- knowledge about the task environment, and the subject's representation of that environment -- have been carefully distinguished and separated. Specifically, SWITCH is given the following information at the start of a simulation run:

- A model of the real world problem (e.g. representations of squares, dominos, and the adjacency relationships between squares)
- A model of the human subject's representation of the real world problem, including concepts that have been generated during problem solving, prior to receiving the PARITY hint (e.g. a concept of a generic square, the proposition that a domino covers two squares)
- An assumed focus of attention (i.e. a 2x2 patch of the board that is referred to first when the simulation needs information about real world squares)
- A set of fairly general productions corresponding to well learned inference rules presumably possessed by adult subjects (e.g. if one proposition appears true based on observation and the same proposition seems false logically, then a contradiction exists).
- Strategic knowledge (implemented in domain specific productions for the purpose of this version of the simulation) corresponding to general strategies such as: "pursue hot ideas" or "change to finer grain size upon failure."
- A hint (corresponding to that given to subjects) that the PARITY (in this case, color) of the squares is important.

How It Works

SWITCH, incorporates three distinct levels of representation. Real World Elements (RWEs) represent the physical problem that exists independent of the problem solver. The Internal Representational Concepts (IRCs) corresponds to the subject's internal representation of these RWEs. Finally, **propositions** are composed of sequences of IRCs strung together. RWEs are necessary in that the simulation must model the task domain, but these elements have no psychological validity. On the other hand, IRCs correspond to representational primitives which subjects combine to form a propositional representation of the problem.

The simulation has two basic capabilities for making progress when stuck: 1) It can try to produce new combinations of the primitives it already has in the hope that these new propositions will trigger some useful knowledge that it has already learned, or 2) it can try to elaborate the IRCs in the hope that changing the primitives themselves will eventually result in useful propositions.

The best way to get a feel for how the simulation works is to examine a production and see what it does. Figure 4 (below) presents (an English translation of) the production that implements the switch from representing "generic" squares to representing squares with color. Note that while the production

may be instantiated with the concepts of "squares" and "color", it corresponds to general knowledge that subjects might have about making analogical mappings.

Production: elaborate-concept-by-analogy

IF: The goal is to prove the problem impossible, AND
 The operator is to elaborate a representation, AND
 A hint exists saying pay attention to some attribute (e.g. COLOR), AND
 Some representational concepts (e.g. the concept of squares) exist
 that have no value for the attribute in question (e.g. COLOR), AND
 There are some real world referents for the representational concept
 that can be referred to (e.g. the squares which really exist
 on the board)

THEN: Map the value of the hinted-at attribute (e.g. COLOR) from
 world objects (e.g. real squares) to the representational
 objects (e.g. representation of squares).

A Sample Production From Switch

FIGURE 4

By analogy to the real world, the simulation is able to shift from an initial representation (IRC) of "square", to a representation (IRC) of "black square" or "white square." A similar production allows the simulation to elaborate old propositions using the new concepts of colored squares. Thus the proposition "A domino covers a square and a square" becomes "A domino covers a black square and a white square."

The transformation of old unelaborated propositions, to new propositions that take parity (color) into account corresponds to the critical shift in representation that allows subjects to make the inferences leading to a Rough Proof. These inferences are modeled in SWITCH by the instantiation and firing of productions corresponding to general knowledge that we believe subjects would possess before beginning the experiment.

The main sequence of SWITCH's behavior in the Checkerboard problem domain, follows:

- 1) Get the hint.
- 2) Decide to elaborate the primitive concepts (IRCs) that form the basis of SWITCH's propositional representation. (Once SWITCH can make no further progress using the stock of propositions it already has, SWITCH follows the strategy of "look at the underlying assumptions".)
- 3) Elaborate primitives (IRCs) by analogy. (The simulation comes up with the new IRCs of "black square" and "white square").
- 4) Decide to generate propositions. (Once new IRCs have been generated, the strategy of "pursue hot ideas" dictates that the simulation check what the implications of the new conceptual primitives will be at the propositional level).
- 5) Elaborate propositions by analogy. (The simulation produces the proposition that "a domino covers a black square and a white square").
- 6) Draw new inferences based on elaborated propositions. (Specifically, SWITCH infers that equal numbers of the two types of squares must be covered, based on the propositions generated in step 5. Pilot data indicates that human subjects make such inferences quite readily.)
- 7) Check inferences against observable fact. (Since the simulation is working within the general context of the "Proof by contradiction" problem space, every new inference should be checked against reality, including the "equal numbers covered" inference made at step 6.)
- 8) Look for contradictions between inferences and observable fact; if found, exclaim "Impossible!"
- 9) Decide to generate a reason for impossibility. (Again, the proof context dictates that the simulation search for a reason for the contradiction.)
- 10) Trace back from contradiction. (The simulation has stored the source of its propositions -- either logically deduced, or empirically observed -- and is able to recall them).
- 11) State Rough Proof. (The simulation uses general knowledge about proofs to frame the information it has recalled).

SWITCH engages in little search and uses only a single problem space in the course of generating a Rough Proof, yet still behaves in a psychologically plausible manner. In contrast, our plans to simulate problem solving behavior prior to the discovery of the concept of parity (Kaplan 1988) call for considerable search in multiple problem spaces.

Without making strong claims for psychological validity at the level of individual productions, we still believe that the overall qualitative behavior of SWITCH is quite similar to that of human subjects who have recognized the significance of parity. Regarding the source of the MC problem's difficulty, the straightforward way in which SWITCH changes representation and then generates a Rough Proof, strongly suggests that the difficulty does not stem from difficulty in reasoning once parity has been noticed. Rather, the problem's difficulty seems to lie in the search that precedes this rapid reasoning.

Human Data on Problem Difficulty

Converging evidence for the proposition that the MC problem's difficulty stems from search comes from analysis of the time spent by subjects at various points along the solution path. Figure 5 illustrates the prototypical solution path along with the mean time spent by subjects at various stages.

After reading the instructions, subjects invariably tried different methods of placing dominos on the board to see if the problem might be solved easily. The length of this covering stage varied widely ranging from about two minutes to almost twenty. It ended when the subject either realized or was told explicitly (via the IMPOSSIBLE and/or INSIGHT hints) that the problem was impossible and that there must be a better approach than trying all possible coverings.

At this point, subjects entered the "search for a new representation" stage, characterized by fewer covering attempts and more search for new approaches to the problem. Some of the common approaches explored in this phase included: symmetry, moving the position of the blue X's, decomposition of the board into smaller boards, various mathematical approaches, counting the number of squares in rows or columns, and finally using the parity of squares (e.g. color). Once subjects focused attention on the parity of squares they were usually able to generate a Rough Proof rapidly.

APPROXIMATE TIME LINE:

SUBJECTS' BEHAVIOR:

MEAN ELAPSED

TIME (n=23 Ss):

	READ INSTRUCTIONS	+++ FORM INITIAL REPRESENTATION
0 Secs. START OF problem.....>		+++
		S
		E
		A
		R
		C
	MAINLY ATTEMPT COVERINGS	H
IMPOSSIBLE HINT (if needed) Given approx. 900 Secs.		S
		E
		A
		R
		C
		H
	SEARCH FOR, & TRY NEW REPRESENTATIONS	S
		E
		A
		R
		C
		H
INSIGHT HINT (if needed) Given approx. 1200 Secs.		S
	FIRST MENTION OF PARITY (COLOR)	E
1194 Secs.>		A
		R
	POSSIBLE ADDITIONAL SEARCH BEFORE RE-FOCUS ON PARITY (COLOR)	C
		H
1508 Secs.....>		+++
	GENERATE Rough Proof	REASONING
1547 Secs.....>		+++
PARITY HINT (if needed) Given approx. 2100 Secs.		
OTHER HINTS (as needed) Given approx. 2450+ Secs.		

A Prototypical Solution Path
FIGURE 5

As Figure 5 indicates, the time subjects spent just *generating* a Rough Proof⁵ is quite small (3%) in comparison with the total time spent problem solving. Most of this proof generation time seemed to be spent finding the right words to communicate the series of rapid inferences that typically followed noticing parity.

The vast majority of subject's time was spent searching -- most of the search (77%) occurring before there is any mention of parity in the protocols. Taken together, the very rapid generation of a proof and the large amount of search before any mention of parity constitute strong evidence that the difficulty of the MC problem stems from search for the correct representation, not reasoning once this representation has been found (as might be the case in other problem solving tasks, e.g. see Wason & Johnson-Laird 1972).

However, it may still seem odd that the gap between first mention of parity and generation of the Rough Proof (approximately 23% of the total time) is as large as it is. Why don't subjects immediately generate a Rough Proof after mentioning parity, if parity really is the key that triggers the new representation as we suggest?

Briefly, the answer involves the fact that the first mention of parity doesn't necessarily indicate a switch of representation or an **understanding** of the importance of parity. In fact, there was considerable variability among subjects in whether noticing parity for the first time led directly to insight. For example, twelve of the twenty-three subjects generated a Rough Proof almost immediately after first mentioning parity, whereas the remaining eleven subjects mentioned parity on two or more separate occasions. Even these eleven latter subjects varied considerably both with regard to the number of mentions of parity and the reasons for its reoccurring mention.

In some cases the initial mention of parity occurred in passing (e.g. "we can cover that... except for the rightmost **pink** one in the second top row") and reflects very little focus of attention. The more interesting (and more common) cases are those in which the subject considers parity but doesn't "see" how it might be of any use. It is likely that many of these subjects were blocked from seeing the relevance of parity by irrelevant knowledge which they refused to put aside. That is, they continued to search for new representations (i.e. problem spaces or approaches) despite an initial (brief) consideration of parity. Even in these cases however, search remains the primary source of problem difficulty. The reason why these subjects seem unable to terminate their search for a representation immediately upon encountering parity will be discussed later in the context of internal sources of search constraint.

External Sources of Search Constraint: Cue Sallience & Hints

If search is the primary source of difficulty, we need to turn to an examination of the potential sources of search constraint in order to understand how subjects manage to change representations and experience insight. We have argued that increasing the salience of parity and providing hints are ways of providing external sources of search constraint.

⁵Time for Rough Proof generation was measured from the moment just before either noticing parity for the first time or last refocussing on parity (in the case of subjects who mentioned parity one or more times but then persisted in exploring other approaches) through the subject's statement of a Rough Proof.

Specifically, we predicted a rank ordering of solution times based upon the expected salience of parity in the four conditions (see bottom of Figure 2). Table 2 presents the relevant results. The first column, showing the mean times required by subjects of different groups to first mention parity, serves as a check on our salience manipulation. As expected, the BREAD & BUTTER board was the most salient (i.e. caused subjects to mention parity earliest), followed by the BLACK & PINK board, the COLOR board, and lastly the BLANK board. A one way analysis of variance indicates that the overall difference between groups is highly statistically significant ($F(3,19)=11.93, p<.0001$).

The second column, confirms our predicted rank ordering for the times to generate Rough Proofs. Again, a one way analysis of variance reveals a statistically significant overall difference ($F(3,19)=4.87, p<.02$).

The third column, showing the number of approaches tried in each experimental condition, offers converging evidence for the same rank ordering suggested by the chronometric results. Some typical approaches were discussed earlier (see Table 1). A one way analysis of variance reveals a highly statistically significant overall difference ($F(3,19)=11.09, p<.002$).

Table 3 indicates which of the differences shown in Table 2 are statistically significant. The BLANK group was statistically different from the other three groups on all three measures. Although the three dependent measures show the same (predicted) rank ordering, differences between the COLOR, BLACK & PINK, and BREAD & BUTTER groups were relatively small, often failing to reach significance. Notice, however, that the groups with the longest solution times also received the most hints (see Table 6 below). Hence the solution times underestimate the differences in performance.

It may also be that problems that require the *invention* of new cues (e.g. the BLANK group) are in a much more difficult class than those that require only *noticing* cues already present (e.g. the remaining three groups). In the former case, without some source of constraints the space of possible inventions is huge, while in case where physical cues are present, there is typically a limited number of features that might be "noticed." In a problem containing relevant cues, subjects need only notice them to constrain their search effectively. We shall discuss later the possibility that one of the distinguishing characteristics of insightful problem solvers is that they are good noticers.

TABLE 2
Mean Times to Solve & 1st Mention Parity, & Mean No. Approaches Tried

Condition	Time to 1st Mention Parity	Time to Rough Proof	Number of Approaches
BLANK	1980 sec.	2273 sec.	9.14
COLOR	1265 sec.	1367 sec.	5.80
BLACK & PINK	905 sec.	1310 sec.	5.16
BREAD & BUTTER	370 sec.	993 sec.	4.00
MEAN TOTAL	1194 sec.	1547 sec.	6.26

TABLE 3
Pairwise Comparisons Between Groups of Ss (df = 19)

Comparisons of Mean Time to 1st Mention Parity:			
	COLOR	BLACK & PINK	BREAD & BUTTER
BLANK	t=2.54, p<.02	t=4.03, p<.001	t=5.73, p<.001
COLOR	----	ns	t=2.95, p<.01
BLACK & PINK	----	----	t=1.84, p<.09

Comparisons of Mean Time to Rough Proof:			
	COLOR	BLACK & PINK	BREAD & BUTTER
BLANK	t=2.32, p<.04	t=2.54, p<.02	t=3.35, p<.004
COLOR	----	ns	ns
BLACK & PINK	----	----	ns

Comparisons of Mean Number of Approaches Tried:			
	COLOR	BLACK & PINK	BREAD & BUTTER
BLANK	t=3.43, p<.003	t=4.30, p<.001	t=5.28, p<.001
COLOR	----	ns	ns
BLACK & PINK	----	----	ns

Hints as Search Constraint

The number of hints required by subjects provides a final source of converging evidence supporting the effectiveness of cue salience as a search constraint. We might expect that subjects whose search is less constrained by elements in the problem (i.e. the low salience subjects) should require more explicit constraint (in the form of hints) in order to solve the problem. However this expectation rests on two important assumptions that have yet to be demonstrated. First we must show that hints are effective at constraining search. Second, we must clarify the relationship between "number of hints required" and "solution time" before we can use the former as a source of converging evidence.

As a test of whether hints actually constrain search, we counted the number of parity statements relevant⁶ to the problem's solution occurring before and after a hint. We expected more relevant statements to occur after hints, thus supporting the hypothesis that hints constrain search.

As the first row of Table 4 shows, subjects did generate more relevant parity statements after hints than before them. Row 2 shows an even greater difference in the number of relevant statements generated before and after the PARITY hint. Both differences are statistically significant (p<.05, one-tailed T test). To check whether these effects might simply result from subjects just talking more after hints, we compared the number of total (i.e. both relevant and irrelevant) statements⁷ generated before and after

⁶The operational definition of a relevant parity statement is identical to the specification of a "relevant invariant" listed in Appendix A. Number of relevant invariants mentioning parity was chosen as the dependent measure since these statements represent the ideal direction in which search should be constrained.

⁷Again Appendix A specifies the coding criteria used here. See relevant and irrelevant invariants.

hints. The mean number generated before hints was 2.08 compared with 2.69 after hints -- far from a significant difference. From these results it appears as if hints not only help, but have their beneficial effect by steering subjects to the appropriate part of the meta-level search space.

TABLE 4
Mean Number of Relevant Parity Statements Before/After Hints

n=8 Subjects			
Type of Hint Considered	No. of hints of this type	Mean no. Statements BEFORE HINT	Mean no. Statements AFTER HINT
Any Hint	13	.46	1.62
PARITY Hint	4	0	3.0

Viewing hints as an external source of search constraint suggests that some hints might constrain search more effectively than others. In particular, since solving the MC problem hinges critically on noticing parity, we might expect the PARITY hint to be particularly effective. Table 5 shows the number of subjects solving the problem after each of the hints. Notice that the PARITY hint appears to have been more effective than all the other hints combined -- at least in terms of the number of subjects who reached a solution soon after a hint was given.

TABLE 5
Number of Subjects Solving After Various Hints
(No. of Subjects in Each Group Who Received Each Hint)

n=23 Subjects					
GROUP	NO HINT	IMPOSSIBLE HINT	INSIGHT HINT	PARITY HINT	OTHER* HINTS
BREAD & BUTTER	3 (5)	0 (2)	0 (2)	0 (0)	2 (2)
BLACK & PINK	2 (6)	1 (3)	1 (3)	1 (2)	1 (1)
COLOR	1 (5)	0 (4)	2 (4)	2 (2)	0 (0)
BLANK	0 (7)	0 (7)	0 (7)	6 (7)	1 (1)
All Groups	6 (23)	1 (16)	3 (16)	9 (11)	4 (4)

*Other hints included the COUNT hint, and other very specific hints which were given only if the PARITY hint proved ineffective.

The numbers in parentheses in Table 5 show how many subjects in each group received a given hint. For example, in the BREAD & BUTTER group, two out of the five subjects received the IMPOSSIBLE hint and the INSIGHT hint (although they did not solve the problem after receiving these hints). Both subjects noticed parity on their own and thus did not require the PARITY hint. However, both subjects were unable to immediately see the connection between parity and the problem's solution (for reasons that will be

discussed in the next section) and eventually had to be told to count the number of different types of squares (the COUNT hint).

It is precisely because subjects can discover the content of a hint on their own (e.g. notice parity, or conclude definitely that the problem is impossible) that the number of hints required is a separate dependent measure from solution time. A subject could take a long time to solve the problem, yet still require relatively few hints. However, if hints act to constrain search, and if subjects in the low salience conditions suffer from a relative lack of external sources of search constraint, we would expect these subjects to require more hints.

The differences in the mean number of hints required ($F[3,19]=2.81$, $p<.07$) are shown in Table 6. The mean number of required hints exhibits the same rank ordering that we have seen in time to 1st mention of parity, solution time, the number of approaches tried, and the percentage of subjects requiring the IMPOSSIBLE, INSIGHT, & PARITY hints (see Table 5). All of these dependent measures converge on the same interpretation: both cue salience and hints are major sources of search constraint which affect subjects' performance on the MC problem.

Table 6
Mean Number of Hints Required

GROUP	MEAN # OF HINTS
BLANK	3.14
COLOR	2.00
BLACK & PINK	1.50*
BREAD & BUTTER	1.20*
ALL GROUPS	2.04
* means significantly different from BLANK group, $p<.01$ (one tailed T test)	

Internal Sources of Search Constraint

While quantitative data such as solution time or the number of hints given provide a good means for understanding the effects of external sources of search constraint on problem solving, such data say little about factors more internal to the problem solver -- factors which may account for individual differences in performance. To explore these factors, we turn to a methodology capable of providing a much richer and denser record of problem solving -- protocol analysis (Ericsson & Simon 1984). Specifically, we shall use protocol analysis to examine the way in which domain knowledge and knowledge of certain general heuristics act to constrain search.

Past research on expertise (e.g. see de Groot 1965, Chase & Simon 1973, Chi, Feltovich, & Glaser 1981) has typically emphasized the power of relevant domain-specific prior knowledge. Less stressed have been the potential adverse effects of bringing such specific knowledge to bear in inappropriate contexts. The first part of our discussion of external sources of search constraint therefore attempts to answer some of the questions raised earlier (e.g. why some subjects solve the problem immediately upon noticing parity while others do not) by examining the adverse effects of misapplying domain specific knowledge.

The second (and major) part of our discussion will focus on heuristic knowledge that may be applicable to a wide variety of problems. After sketching a sample of the heuristics used by subjects to constrain their search, we will focus on the Notice Invariants heuristic as a potential source of individual differences. We will see that this heuristic separates cleanly the good from the poor problem solvers, regardless of experimental condition, suggesting that generality of a heuristic need not necessarily be bought at the price of lessening its power.

Prior Knowledge as Search Constraint

Knowledge is a two edged sword. For most problems, knowledge allows one to hack away irrelevant details and focus on the problem elements that are likely to be critical for a solution. But in the case of insight problems, where the answer often lies in a very obscure place, knowledge can cut the other way. Inappropriate or irrelevant knowledge guides search to an unproductive region of the problem space, as we saw earlier in the case of the unfortunate Chemical Engineering graduate student.⁸

As an illustration of the way in which knowledge constrains search, consider the case of S9, a subject in the BREAD & BUTTER condition. S9 was a sophomore majoring in Chemistry who did not consider himself to be a puzzle solver. His behavior is of particular interest because he took a long time to solve the problem despite the facts that he was in the easiest condition and that he noticed parity without the benefit of the PARITY hint.

S9 began by reading the instructions, calculating whether the *number* of dominos was sufficient to cover the number of remaining squares, and trying a number of coverings:

So, I have 31 dominos, they cover two squares apiece, that's 62 squares ...
...logically, it should cover it...
...maybe we can cut out this section of the board right here. This middle
section, that gives 6 x 6 that's ... 36 ... 36 squares... and that can be ...
covered.

Notice that even in these beginning statements, S9's behavior is far from random search. His calculations provide constraints indicating that a COVERINGS problem space is likely to contain the solution. Specifically, here he conceptualizes the problem in terms of cutting out sections of the board rather than placing individual dominos.

As his covering attempts fail, S9 searches for invariants that might serve as explanations. Such reasons might also act to constrain and direct his search.

... we seem to missing one [domino placement] always ...
... we still seem to be missing one ...
... Argh [pause] missed it by one ...
... this an eight by eight checkerboard [pause] and it should be able to
cover [pause] but since both of these numbers are [pause] even, and 31 is odd
[pause] but it does make 64. I'm gonna say that [pause] it can't be done.

Although unconvinced of the reason, S9 has detected an invariant – namely that all his attempts seem to fail. This invariant argues that the problem is impossible, while his early calculations indicate that a covering should exist. Lacking a clear way to resolve the conflict, S9 persists both in trying coverings and

⁸Other studies have found the similar results in the other domains, for example the domain of algebra word problems (Paige & Simon 1966).

in stating periodically that the problem is impossible. It is interesting that other subjects are able to give up their belief in the problem's possibility much more easily than S9, who seems to need a reason to explain the apparent contradiction. After fifteen minutes has elapsed, the experimenter confirms that the problem is indeed impossible and that the task is to find a logical proof of this fact. Upon receiving this hint, S9 generates a series of potential explanations for the problem's impossibility:

... the two Xs have to be side by side in order for it to be done ...
 ... Maybe cause it's [pause] it's 31 which is an odd number, maybe? uh
 [pause] even though it is multiplied by two [pause] it [pause] doesn't go into 64? ...
 ... I think it's because [pause] ~~unmmmm~~ if you were able to cover it with 32 dominos, and have no Xs, then 32 is like 2 to the 5th [pause] and ~~unmm~~ [pause] 64 is 2 to the 6th. If [pause] it's only off by a factor of 2 and it's multiplied by a factor of 2 [pause] and [pause] I said this before I guess ... if these two Xs are not side by side -- with the exception of diagonally -- then you leave two open spaces far away from each other ...

Here we see the influence of S9's domain specific mathematical knowledge on what he notices. S9 is noticing invariant properties of the board, but he focuses on mathematical invariants rather than on invariants having to do with perceptual features (e.g the words Bread and Butter). However, it is not the case that S9 simply fails to notice Bread & Butter. Early during problem solving, he mentions:

This is tough. Come on [pause] bread and butter [pause] 64, two squares [pause] it should cover [pause] uhhhh could fit in a logical pattern.

And later:

... could it be bread will not fit on butter? [long pause] can't find [pause] a logical way ...

S9 seems to have two competing sources of search constraint available, the highly salient cues, Bread & Butter, and his knowledge of mathematics which seems relevant to this problem. His insistence that the solution "be logical" reflects his mathematical approach. Thus, although S9 returns to Bread & Butter repeatedly as his mathematical approaches fail, he is extremely reluctant to explore this apparently illogical approach. Eventually, however, the continued failure of his mathematical ideas drives him to ponder:

Why is the bread and butter on the [pause] Ahh. OK. Like there are two on [pause] there's two that you need to have butter on [pause] Ohhhh [pause] that leads to something. If we say butter is [pause] if [pause] if we have a perfect situation [pause] we [pause] in a perfect situation, you see there's butter on one and then bread on the other. But since [pause] if the number of spaces [pause] like when the Xs are apart, is odd, we end up with butter on both things ...

At this point, S9 notices that the parity of the squares co-varies with the even/odd number of spaces between the Xs. That is, if the board were coverable (a perfect situation) then one X would be on a bread, the other on a butter, and the number of spaces between the Xs would be even. An uncoverable board however might have both Xs on butters and an odd number of spaces between the Xs.⁹

⁹The reader can verify that even/odd spacing is a true invariant in the following manner: Turn to the checkerboard illustrated in Figure 1. Starting with any square next to one of the Xs, find a path to any square next to the diagonally opposite X, using only horizontal or vertical movements. No matter what the shape of the path, the number of squares it contains will prove to be odd. However, if one of the Xs is moved so that the problem is possible (i.e. each X is on a different color), and a new path is found by the same method, now the path contains an even number of squares.

Typically, subjects who have noticed the correspondence between the parity of the removed squares and possibility of the problem are able rapidly to generate a Rough Proof. For S9, the only critical fact that is missing appears to be the notion that a domino must cover one square of each type. We would expect that all the attention being focused on bread and butter should allow S9 to note this readily available fact. Not so.

S9 remains unwilling to explore his new discovery without a clear idea of its logical relevance. We noted earlier that S9 was reluctant to relinquish the idea that the problem was possible, although he noted the failure evidence quite early on. Similarly, S9 remarked on the existence of BREAD & BUTTER before he even decided that the problem was impossible, yet here at this late stage in problem solving he is still reluctant to explore its ramifications. Instead he persists in trying to relate the parity invariants to his previous mathematical ideas:

...if you designate butter odd [pause] as odd [pause] and bread as even [pause] then in many -- any -- situation where we're trying to figure out [pause] a square root of a number minus [pause] two [pause] we need odd even...
 ... But logically this doesn't make sense [pause] I can't explain this logically. OK, but it works logically, [pause] this has no cause and effect...

Ultimately S9 must be pried loose from his mathematical preconceptions with the hint to count the numbers of bread and butters on the board. Only after this hint does S9 realize that a domino covers one of each type. Then he is able to solve the problem.

The behavior of S9 is in direct contrast to the behavior of another BREAD & BUTTER subject who solved the problem quite rapidly, S1. Like S9, S1 noticed the words BREAD & BUTTER early in problem solving. Unlike S9, however she had no reservation about exploring their potential relevance. In fact, she startled the experimenter by reading the problem instructions and immediately asking:

The ... the words don't matter do they? ... in this problem?

The experimenter answered, "No," believing that any other response would give away the answer before the problem was even begun. Since she was put on the wrong track at the start of the problem, one might expect that S1 would require a long time to reach a solution. Not so. While S9 had been overly concerned with mathematical rigor and logical connections, S1 rapidly suggests one "proof" after another -- most of which amount to statements that particular covering(s) won't work. When these proofs are rejected, she returns to Bread & Butter (see Figure 3 for an excerpt from her protocol), and solves the problem quite rapidly.

Inexperience with the nature of formal proofs (as evidenced by her attempts to call failed coverings "proofs") together with her willingness to try new approaches worked to S1's advantage. While S9's understanding of proofs, and his knowledge of mathematics, apparently led him to concentrate on trying to solve the MC problem mathematically, it was the apparent lack of such (proof) knowledge that may have helped S1 achieve her rapid solution.

Thus, the availability of domain specific knowledge provides at least one answer to the question of why some subjects solve the problem immediately after noticing parity, but others do not. To explore other reasons why some subjects are fast solvers while others wander down long and hopeless paths, we must

turn to the final source of search constraint, heuristics.

Heuristics

So far we have examined cues in the problem, hints, and domain specific prior knowledge (e.g. mathematics) as sources that subjects might use to constrain their search in a vast space of potential representations. If we were to stop our investigation at this point, what could we say about the process of insight?

Given time to analyze a particular insight problem (and its solution), we might be able to predict how a subject with specified prior knowledge would perform (at least qualitatively). We could probably also suggest how modifications to the task would make the problem easier or more difficult (based on the cue salience results). Finally we might be able to devise hints, and use them with a better understanding of how they have their effect.

But what might we do to promote insights in domains where we did not already know the answer? What leads one person to insight, and not another?

We believe that the answer to these questions lies partly in the heuristics people use. The first step in exploring this hypothesis is to get some hard evidence that people use heuristics in general. While this fact has been more or less assumed so far (and in fact appears evident in the protocol of S9) we counted the usage of some general heuristics to put the assumption on firm ground.

General Use of Heuristics

We hypothesized that subjects might use at least three very general heuristics to help constrain their search: Noticing Invariants, Forming Hypotheses, and Comparing Alternative Board Situations. The notes made from the 23 verbal protocols reveal that each subject used each of these heuristics at least once. A more detailed analysis of the eight verbatim transcripts provided the data shown in Table 7.

The clearest result is that subjects used the Notice Invariant heuristic more often than Hypothesize, which in turn was used more often than Compare. This ordering is intuitive since we would expect subjects to notice more facts than they actually hypothesize about, and to form more hypotheses than they actually test.

TABLE 7
Overall Frequency of Use (and rate of use) of Heuristics
n=8 Subjects

SUBJECT	NOTICE-INVARIANTS	HYPOTHESIZE	COMPARE
1 BREAD & BUTTER (Fast)	26 (2.7/min.)	7 (.73/min.)	3 (.31/min.)
14 BLACK & PINK (Fast)	9 (1.7/min.)	6 (1.1/min.)	1 (.19/min.)
2 Color (Fast)	27 (2.6/min.)	3 (.29/min.)	2 (.20/min.)
17 BLANK (Fast)	33 (1.0/min.)	20 (.61/min.)	5 (.15/min.)
9 BREAD & BUTTER (Slow)	47 (1.1/min.)	29 (.70/min.)	8 (.19/min.)
13 BLACK & PINK (Slow)	43 (1.1/min.)	32 (.85/min.)	6 (.16/min.)
16 Color (Slow)	31 (1.4/min.)	20 (.91/min.)	5 (.22/min.)
4 BLANK (Slow)	107 (2.3/min.)	51 (1.1/min.)	12 (.25/min.)
TOTAL	323	168	42
MEAN	40.4	21.0	5.3

Noticing Invariants

In our attempt to describe the processes underlying insight, we have already described a means for accomplishing the actual switch of representation (i.e. the SWITCH simulation discussed above). What remains is the question of what leads subjects to consider the critical cues which are prerequisite for such a switch.

The Notice Invariants heuristic is of central importance to our account of how subjects, particularly those not given strong perceptual or verbal hints, are able to narrow the space they search for new representations. The space of "all possible representations" is ill-defined and certainly unmanageable. Subjects are not equipped with generators for searching a space like that. How do they, in the MC problem, generate representations in which the relevant feature, the color of the squares, plays a central role?

We hypothesize that, while solving a difficult problem, people are attentive, at least intermittently, to features of the problem display. In particular, if some features are invariant -- do not change as the situation changes -- they will sooner or later attract attention and be remembered. For example, in the MC problem, the color of the two squares that cannot be covered in a covering attempt are always of the same color, and that color is the opposite of the color of the two squares that have been removed. Another invariant (this one not directly relevant to the solution) is that the problem is solvable if the two squares removed are an even number of squares apart, insoluble if they are an odd number of squares apart.

Table 8 lists some common invariants noticed by the eight subjects. All the invariants relevant¹ (i.e. on the direct solution path) to the insightful solution are listed as well as the more common irrelevant invariants (i.e. those invariants that are not on the direct solution path). The range of invariants noticed

¹See Appendix A for a precise definition of "relevance" and "invariant"

demonstrates the generality of the heuristic. Precisely because Noticing Invariants is a widely applicable rule of thumb for searching in ill-defined domains, there can be no guarantee that the invariants noticed will be the critical ones for the particular problem. Nevertheless, the constraints offered by the Notice Invariant heuristic are a vast improvement over blind trial and error search. Specifically, noticing invariants provides a generator for possible problem spaces, a large fraction of which incorporate relevant invariants.

TABLE 8
Invariant Properties Mentioned Repeatedly by Subjects
 n=8 Subjects

Description of Invariant Property	% Ss (n=8) who noticed it	Relevance
Problem is impossible /Coverings Fail	88%	Relevant
Domino covers two of different parity	75%	Relevant
Restatement of the givens (e.g. 31 dominos 62 squares left) or the goal	63%	Relevant
Squares of the same parity are removed/two removed of the same parity --> impossible	50%	Relevant
Adjacent squares are of different parity/ diagonal squares are same parity	38%	Relevant
A domino covers two adjacent squares/ domino cannot cover diagona'ly	25%	Relevant
An imbalance exists between the number of squares of different parity	13%	Relevant
(Two diagonal) squares are always left over when a covering attempt fails.	63%	Irrelevant
Various mathematical properties related to the (number of) squares or dominos	63%	Irrelevant
Various patterns of covering/symmetries of the board	50%	Irrelevant
The problem's possibility depends upon the position (not parity) of the Xs	50%	Irrelevant
Odd # of spaces between Xs indicates problem impossible, even # --> possible	38%	Irrelevant

Note: Listed first are all the invariants on the direct path to the insightful parity solution (i.e. the relevant invariants). Listed second are some common invariants that are NOT on the direct path to the insightful parity solution (i.e. the irrelevant invariants).

Individual Differences


There are at least two potential sources of individual differences suggested by our focus on the Notice Invariant heuristic. First, there may be differences in the number of things noticed. Our measure here would be the total number of invariants generated. Second, there may be differences in the types of things noticed. In this case we would want to break the invariants according to the categories described in Appendix A (e.g. relevant, irrelevant, etc.).

Quantity of Invariants

We have already seen from Table 7 that there is no significant difference in the overall rate of generating invariants between fast and slow subjects. However a more detailed analysis requires that we examine the number of invariants generated at different times during problem solving. When the eight protocols are broken up into time slices (each roughly 100 seconds long), we find that the fast subjects generate significantly more invariants than the slow subjects in the first five minutes ($p < .05$ one tail t test). Fast subjects appear to notice more things, earlier.

Figure 6 illustrates the differences between fast and slow subjects in the mean cumulative number of invariants generated. Since one fast subject solves the problem after approximately 300 seconds, another after approximately 500 seconds, and a third after approximately 600 seconds, the means for the fast subjects are based on progressively fewer data points (explaining the drop in mean cumulative frequency from time slice 5 to 6). However, it is clear that the fast subjects take a quick lead in the number of invariants generated and that the slow subjects are unable to catch up until after some of the fast subjects have already solved the problem (and therefore have ceased to generate new invariants).

This result suggests that fast subjects are those who constrain their search the most by rapidly generating invariants from the start. Slow subjects might eventually "catch up", but not until most of the fast subjects (the notable exception being the fast subject in the BLANK condition where the relevant cue is unavailable) have already noticed important facts about the problem and solved it!



Cumulative Frequency of All Invariants (1st 10 min.)

FIGURE 6

Perceptual Invariants


While there are clear differences in the number of things noticed by fast and slow subjects during the first 10 minutes, we might wonder: Will noticing any of fact at all help, or does noticing one type of invariant rather than another play a key role?

Both the cue salience results and the phenomena of interference from prior knowledge suggest that paying attention to the unique properties of the MC problem is likely to meet with success. In the case of the MC problem, the unique properties tend to be perceptual properties of the board. The instructions are quite simple, so subjects who tend to look to the problem for constraints on search end up literally staring at the board, and noticing. Could this be what the fast subjects are doing?

To test this hypothesis we categorized invariants according to whether they were perceptual (e.g. related to color, the position of Xs, or other visual aspects of the problem) or non-perceptual (e.g. related to strategies such as decomposition, to mathematical approaches, or to other conceptual rather than visual properties of the problem).


Figures 7 & 8 shows the cumulative frequency of non-perceptual invariants and perceptual invariants generated by fast and slow subjects during the first 10 minutes of problem solving. Figure 7 shows that the pattern for non-perceptual invariants is quite similar to the pattern for invariants in general (Figure 6). Fast subjects seem to generate more in the first five minutes; then the slow subjects catch up. Again, we believe this reflects a general tendency for the fast subjects to seek sources of search constraint immediately by noticing things from the very start of problem solving.

The results in Figure 8 are much more striking. Here we find a clear separation between fast and slow subjects that extends over the first ten minutes. Every single fast subject generated one or more perceptual invariants within the first ten minutes, while none of the slow subjects did so! In fact, although the graph does not show this, the slow subjects *never* catch up with the fast subjects on this measure until after all of the fast subjects have already solved the problem.



Cumulative Frequency of Non-Perceptual Invariants (1st 10 min.)

FIGURE 7



Cumulative Frequency of Perceptual Invariants (1st 10 min.)

FIGURE 8


Eliminating an Uninteresting Explanation

One uninteresting explanation of these results would be that fast subjects (almost by definition) notice more relevant things about the problem. It could be that the difference shown in Figure 8 simply reflects the fact that fast subjects are noticing relevant things, whereas slow subjects are not. In response to this argument, we point out that not all of the perceptual invariants are necessarily relevant to the problem solution (e.g. noticing that the Xs are diagonal from each other, while a perceptual property, is not on a direct path to the solution). Conversely, non-perceptual properties (e.g. the knowledge that a domino is only allowed to cover adjacent squares) can also be relevant. A more convincing argument, however is made by the data on relevant invariants plotted in Figure 9.

Here we see that the pattern is quite different from Figure 8 in two respects. First, slow subjects generate **relevant** invariants from the very start of the problem, in contrast to their striking failure (shown in Figure 8) to generate **perceptual** invariants. Second, slow subjects start off slightly ahead of the fast subjects, rather than being behind from the start. Therefore, being "fast" does not translate to simply "generating more relevant invariants."

The increase in the number of relevant invariants generated over time (for both groups) is what we would expect if subjects are "homing in" on the solution over time. The rather sharper increase on the part of the fast subjects is presumably due to the facts that more and more relevant invariants are generated as subjects approach the solution, and that three out of the four fast subjects solve the problem within the first ten minutes.


Converging evidence for this "homing in" interpretation comes from counting the number of relevant and irrelevant invariants generated by subjects in the first and second half of their protocols. Figure 10 shows that subjects notice more total invariants in the first half of their protocols than in the second half. However, as the "homing in" hypothesis predicts, subjects notice significantly more **relevant** invariants ($p < .05$ one tailed) and significantly less irrelevant invariants ($p < .05$) in the second half of their protocols.



Cumulative Frequency of Relevant Invariants (1st 10 min.)

FIGURE 9

n=8



Number & Type of Invariants Noticed in 1st & 2nd Halves of Protocol

FIGURE 10

Having eliminated the most obvious alternative explanation, we are left with the interpretation that fast subjects possess a heuristic that the slow subjects don't -- namely that of paying attention to the perceptual features of the problem. It may even be that the sooner subjects use this heuristic, the sooner they solve the problem.

To test this more specific hypothesis, we could compare the order in which subjects mentioned their first perceptual invariant with the order in which they solved the problem. Table 9 makes just such a comparison. The rank order correlation between first mention and solution is .98 for the fast subjects, and .86 overall. At least for fast subjects, it seems that noticing perceptual invariants is almost a perfect predictor of speed in solving the MC problem. As you might expect, the difference in times between fast and slow subjects (in the same experimental conditions) to first mention of perceptual invariants is highly significant ($p < .01$ one tailed T test).

TABLE 9
Rank Orders of Times to 1st Mention Perceptual Invariant & to Solution
n=8 Subjects

Subject & Condition	Rank Order Time to 1st percept. Invar.	Rank Order Time to problem Solved
s14 P&B fast	1	1
S1 B&B fast	2	2
S2 Clr fast	3	3
S17 Blnk fast	4	5
S13 P&B slow	5	6
S9 B&B slow	6	7
S16 Clr slow	7	4
S4 Blnk slow	8	8


Both the differences in the number of perceptual invariants mentioned, and the correlation data above strongly support the view that fast subjects are using qualitatively different heuristics than slow subjects. The exact nature of these heuristics is open to debate. Because of the Checkerboard problem's perceptual nature, one could argue that fast subjects are simply perceptually driven and that they would not be able to apply the more general heuristic of "pay attention to problem features" in other domains. This is, of course, an empirical question. However the more general result that fast subjects generate more invariants overall, suggests that a tendency to notice visual features in and of itself is not responsible for the performance of the fast subjects. From their own experience, most researchers can certainly attest to the benefit of "listening to what the data is saying" -- a task that does not always involve visual features.

Flexibility In Noticing

As a final exploration of individual differences, we tried to test the relation between the age old advice of "Think flexibly!" and the promising heuristic of "Pay attention to invariants!". Assuming that flexibility means generating ideas of different types, one logical measure of flexibility would be to count the number of different types of invariants (where type is defined by the coding categories listed in Appendix A)

generated by fast and slow subjects. Since the first fast subject solved the MC problem at the end of five minutes, we can only present a full comparison for the first five minutes. Figure 11 shows the mean number of coding categories (for invariants) covered by fast and slow subjects during this time interval.

The difference between the number of categories covered by fast and slow subjects is statistically significant ($p < .02$ one tailed). In one respect this result is not too surprising since we already know that fast subjects generate more invariants overall. However, since fast subjects notice significantly more non-perceptual invariants than slow subjects during this same time period (Figure 7), the difference in category coverage cannot be due primarily to the difference in the number of perceptual invariants noticed. Rather it shows that fast subjects are noticing not only more things, but a wider variety of things. Future experiments may help determine whether it the absolute number of invariants noticed, the breadth of noticing, or both, which constrain search most effectively.



Categories Covered (1st 5 min.)

FIGURE 11

Conclusions

We began our story with a metaphor comparing solving insight problems to searching for a diamond in a dark room. We argued that the task would be hopeless without some source of search constraint. In the domain of the MC problem, we identified four potential sources of search constraint: Perceptual cues in the problem, hints provided by the experimenter, prior knowledge that the subject might bring to the problem, and heuristics -- in particular the heuristic of Notice Invariants.

After establishing that the difficulty of the MC problem stems from search, we examined each of the potential sources of search constraint in turn. The fact that subjects solving more salient versions of the MC problem attained insight sooner attests to the power of cue salience as a source of search constraint. Hints, especially the PARITY hint, were shown to be quite effective sources of search constraint. Subjects generated more statements relevant to the solution path after a hint than before it. Protocol evidence suggested that prior domain knowledge, while constraining search, could actually be counterproductive if it leads to search in the wrong part of the space (as it is likely to do with insight problems).

Perhaps our most interesting results concerned use of the Notice Invariant heuristic. Focusing attention on invariant features of the problem situation guarantees that the ideas considered have at least a minimal relevance to the problem. Therefore it allows subjects to convert a search in an enormous and unmanageable space (in which they have no relevant generators) to a search in a small space (with generators available).

We found that all subjects used this heuristic commonly, and that fast subjects used it more often than slow subjects early in problem solving. Furthermore, fast subjects differed from slow subjects in the types of invariants they noticed. In particular, of the eight subjects whose behavior was examined in detail, all fast subjects noticed perceptual invariants within the first ten minutes of problem solving, whereas none of the slow subjects did so. This result cuts across experimental conditions, clearly separating the fast from the slow. Fast subjects also tended to notice a wider variety of invariants than slow subjects during the initial minutes of problem solving, suggesting that flexibility, or the willingness to try a variety of things may facilitate insight.

The "Notice Invariant" results constitute a significant step towards identifying a heuristic that can facilitate insight across a wide variety of domains. Although hints, cue salience, and prior knowledge all constrain search, it is difficult to specify how they will have their effect without knowing the nature of the problem (and its solution) beforehand. The essence of discovery however, is that you do not know beforehand where the solution may lie. If noticing invariants, and in particular perceptual invariants, provides even a little search constraint for the ill-defined task of discovery, then we have a cause for celebration. As this paper has tried to show, a little search constraint goes a long way.

References

- Anderson, J.R. (1985) *Cognitive psychology and its implications*. 2nd. edition. New York: W.H. Freeman & Co.
- Burnham & Davis (1969) The 9-dot problem: Beyond perceptual organization. *Psychonomic Science*, 17, 321-323.
- Chase, W.G. & Simon, H.A. (1973) The mind's eye in chess. In W.G. Chase (Ed.), *Visual information processing*. New York: Academic Press.
- Chi, M.T.H., Feltovich, P.J., & Glaser, R.J. (1981) Categorization and representation of physics problems by experts and novices. *Cognitive Science*, 5, 121-152.
- Davidson, J.E. (1986). The role of insight in giftedness. In R.J. Sternberg & J.E. Davidson (Eds.), *Conceptions of giftedness*. New York: Cambridge University Press.
- de Groot, A.D. (1965) *Thought and choice in chess*. The Hague: Mouton.
- Duncker, K. (1945) On problem solving. *Psychological Monographs*, 58, (Whole No. 270).
- Durkin, H.E. (1937) Trial and error, gradual analysis, and sudden reorganization: An experimental study of problem solving, *Archives of Psychology*, no. 210.
- Ericsson, K.A. & Simon, H.A. (1984) *Protocol analysis*. Cambridge, MA: MIT Press.
- Ghiselin, B. (1952) *The creative process*. Berkeley, CA: University of California Press
- Haefele, J.W. (1962) *Creativity and innovation*. New York: Reinhold Publishing Corp.
- Hadamard, J. (1949) *The psychology of invention in the mathematical field*. Princeton, N.J.: Princeton University Press.
- Kaplan, C.A. (1988) SWITCH: Towards a simulation of representational change in the Mutilated Checkerboard Problem. *Technical Report*, Carnegie-Mellon University.
- Katona, G. (1940) *Organizing and memorizing*. New York: Columbia University Press.
- Kohler, W. (1956) *The mentality of apes*. London: Routledge & Kegan Paul Ltd.
- Kohler, W. (1969) *The task of gestalt psychology*. Princeton, N.J.: Princeton University Press.
- Korf, R.E. (1980) Toward a model of representational changes. *Artificial Intelligence*, 14, 41-78.

- Laird, J.E., Newell, A., & Rosenbloom, P.S. (1987) Soar: An architecture for general intelligence. *Artificial Intelligence*, 33, 1, 1-64.
- Langley, P., Simon, H.A., Bradshaw, G.L., & Zytkow, J.M. (1987) *Scientific discovery: Computational explorations of the creative process*. Cambridge, MA: MIT Press
- Larkin, J.H., Reif, F., Carbonell, J., & Gugliotta, A. (1985) Fermi: A flexible expert reasoner with multi domain inferencing. *Technical Report*, Department of Psychology, Carnegie-Mellon University.
- Lung, C. & Dominowski, R.L. (1985) Effects of strategy instructions and practice on nine dot problem solving. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 11, 4, 804-811.
- Maier, N.R.F. (1931) Reasoning in humans: II. The solution of a problem and its appearance in consciousness. *Journal of Comparative Psychology*, 12, 181-194.
- McCarthy, J. (1964) A tough nut for proof procedures, *Stanford Artificial Intelligence Project Memo* 16, July 17.
- Newell, A. (1965) Limitations of the current stock of ideas about problem solving. in A. Kent & O.E. Taulbee (Eds.) *Electronic information handling*. Washington: Spartan Books.
- Newell, A. & Simon, H.A. (1972) *Human problem solving*. Englewood Cliffs, N.J: Prentice-Hall
- Newell, A., Shaw, J.C., & Simon, H.A. (1962) The process of creative thinking. in H.E Gruber, G. Terrell, & M. Wertheimer (Eds.) *Contemporary approaches to creative thinking*. New York: Atherton
- Paige, J.M. and Simon, H.A. (1966) *Cognitive processes in solving algebra word problems* reprinted in H.A. Simon *Models of thought*. New Haven: Yale University Press, 1979.
- Posner, M.I. (1973) *Cognition: and introduction* Glenview, ILL: Scott, Foresman & Co.
- Simon, H.A., & Lea, G. (1974) Problem solving and rule induction. reprinted in H.A. Simon *Models of thought*. New Haven: Yale University Press, 1979.
- Wason, P.C. & Johnson-Laird, P.N. (1972) *Psychology of reasoning: Structure and content*. Cambridge, MA: Harvard University Press.
- Weisberg, R.W. & Alba, J.W. (1981) An examination of the alleged role of "fixation" in the solution of several "insight" problems. *Journal of Experimental Psychology: General*, 110, 2, 169-192.
- Wertheimer, M. (1945) *Productive thinking*. New York: Harper & Row
- Wickelgren, W.A. (1974) *How to solve problems*. San Francisco: Freeman.
- Worthy, M. (1975) *Aha! A puzzle approach to creative thinking*. Chicago: Nelson Hall.

APPENDIX A: CODING SYSTEM

Coding of the protocols occurred in three phases. In the first phase, the tapes were transcribed verbatim. Eight of these verbatim transcripts were made -- one for each of the eight subjects selected for detailed analysis. We listened to the tapes of the other 15 subjects and made notes about their content and the timing of critical events (e.g. when color was first mentioned, when a rough proof was generated, etc.). However, these remaining 15 tapes were not transcribed verbatim.

Each of the 8 verbatim transcripts were coded for occurrences of INVARIANTS, COMPARISONS, and HYPOTHESES in the second phase.

These categories were defined as follows:

INVARIANTS:

An invariant is a fact that is mentioned repeatedly, and/or is qualified with one of the words "always", "any", "every", or "never."¹¹

COMPARISON:

A comparison is the mention of two actual or hypothetical board situations in a single sentence, or in two consecutive sentences.

HYPOTHESIS:

A hypothesis is defined as an "if" statement that proposes, or refers, to a situation that could exist or an action that could be taken.

These categories are not mutually exclusive. For example the phrase "if the covering always fails ..." would be coded both as an invariant and a hypothesis. Invariants, hypotheses, and comparison were coded on separate passes through the transcripts using fresh copies of the transcript for each pass. The phrase or group of words matching the category was marked with a highlighter pen.

The third phase of coding involved categorizing the content of each INVARIANT as one of nineteen mutually exclusive types. The content of each INVARIANT was matched against the defining content (information in parentheses below) of each type in the following order:¹²

- 1) Color-Imbalance (incl.: unequal #s of two colors, 32 black & 30 pink or vice versa, 2 pinks or blacks MUST be left)
- 2) Color-Covered (incl.: domino covers a pink & black, domino covers 2 of different colors)
- 3) Color-Position (incl.: pink & black adjacent, squares of same color are diagonal, pink & black go together, pink are left)
- 4) Color-Removed (incl.: same color squares Xed, Xs are both black, Xs on pink & black --> possible, Xs on same color --> impossible,)

¹¹These criteria were used to filter out facts that a subject might state in passing. We wished to identify those facts that subjects believed to be important and invariant. In fact, subjects' notion of what facts were invariant coincide quite well with objective analysis.

¹²Each INVARIANT was categorized based upon the first successful match with a type.

- 5) Infer-Possibility (incl.: the problem is impossible)
- 6) Infer-Coverability (incl.: domino covers two adjacent squares, 2 remaining squares are not coverable, one domino remains after covering 30 squares)
- 7) Infer-Position (incl.: horizontal or vertical squares are adjacent)
- 8) Given-Covering-Properties (incl.: domino covers horizontally or vertically, domino cannot cover diagonally)
- 9) Given-Resources (incl. 8x8 board, 64 squares, 62 squares remaining, 2 sq. removed, 31 dominos available)
- 10) Given-Goal (incl.: prove logically impossible, find covering)
- 11) Other-Color (incl.: any statements mentioning the COLOR of squares, excepting statements specifically covered above)
- 12) Xposition & Possibility (incl.: move Xs, Xs adjacent --> possible, Xs moved --> possible, even/odd spacing & possibility, Xs not diagonal --> possible, position of Xs is responsible for impossibility, Xs make problem impossible)
- 13) Xposition-General (incl.: Xs are diagonal, Xs exist)
- 14) Type-Left (incl.: two squares always left, squares left are diagonal)
- 15) Covering-Failures (only 30, not 31, can't do it, doesn't work, etc.
EXCLUDES DIRECT STATEMENTS OF PROBLEMS IMPOSSIBILITY)
- 16) Math (incl.: even/odd numbers --NOT SPACING --, algebra, facts about the numbers involved, counting features of board)
- 17) Decomposition (incl.: CONSIDERING nxn boards, mentioning decomposition, trying different shaped boards)
- 18) Covering (incl. ALL STATEMENTS ABOUT SPECIFIC DOMINO PLACEMENTS AND STRATEGIES OR SYSTEMATIC DOMINO PLACEMENT)
- 19) Other (incl.: Off-track ideas of rare frequency)

Once all INVARIANTS had been coded according to content, it was possible to define subsets of the nineteen content types that define other meaningful categories. Specifically:

Types 1 - 10 = RELEVANT INVARIANTS, i.e. invariants that we would expect to be generated by a subject following a direct path to the insightful solution.

Types 11-19 = IRRELEVANT INVARIANTS, i.e. invariants that do not directly lead to the insightful solution.

Types 2-4, 11, 13, 14 = PERCEPTUAL INVARIANTS, i.e. features/properties of the board that are evident at a glance.

Types 1, 5-10, 12, 15-19 = NON-PERCEPTUAL INVARIANTS

Types 1-4 = RELEVANT COLOR INVARIANTS.

Type 18 = COVERING INVARIANTS

APPENDIX B: Stimulus Materials

Instructions for Subjects in the Bread & Butter Experiment

INSTRUCTIONS

Imagine that you have 31 dominos. Each domino is big enough to cover exactly two squares if it is placed horizontally or vertically. The domino cannot cover two diagonal squares however.

The experimenter will cover two of the squares on the board with blue Xs. Your problem is to decide if it is possible to cover all the remaining squares (i.e. all those except the ones covered with blue Xs) with the 31 dominos. If you think it is possible, you must show how you would do it. If you think it is impossible, you must PROVE logically why the problem is impossible.

If you have any general questions, you may ask before you begin. Remember to THINK ALOUD as you puzzle.

The Bread & Butter Board

Subject _____

bread	butter	bread	butter	bread	butter	bread	butter
butter	bread	butter	bread	butter	bread	butter	bread
bread	butter	bread	butter	bread	butter	bread	butter
butter	bread	butter	bread	butter	bread	butter	bread
bread	butter	bread	butter	bread	butter	bread	butter
butter	bread	butter	bread	butter	bread	butter	bread
bread	butter	bread	butter	bread	butter	bread	butter
butter	bread	butter	bread	butter	bread	butter	bread


The Pink & Black Board

Subject _____

black	pink	black	pink	black	pink	black	pink
pink	black	pink	black	pink	black	pink	black
black	pink	black	pink	black	pink	black	pink
pink	black	pink	black	pink	black	pink	black
black	pink	black	pink	black	pink	black	pink
pink	black	pink	black	pink	black	pink	black
black	pink	black	pink	black	pink	black	pink
pink	black	pink	black	pink	black	pink	black


The Color Board

Subject _____



The Blank Board

Subject _____



Experimenter's Data Sheet**SUBJECT INFORMATION & COMMENTS FORM****SUBJECT #** _____**Date** _____**EXPERIMENTAL CONDITION:** _____
_____**MAJOR:** _____**Year** _____**Previous Psych. classes:** YES NO**SEX:** M F**Previous Problem Solving classes:** YES NO**Is subject "a puzzler"?** _____**Favorite type of puzzle:** _____**Last puzzle done:** _____**Says has seen Monk problem:** YES NO

+++++

In Search of Insight

Craig A. Kaplan and Herbert A. Simon
Carnegie-Mellon University
15 August 1988

This research was supported in part by the Defense Advanced Research Projects Agency under Contract F-33615-84-K-1520, and in part by the Personnel and Training Programs, Psychological Sciences Division, Office of Naval Research, under Contract No. N00014-86-K-0768. Reproduction in whole or in part is permitted for any purpose of the United States Government. Approved for public release; distribution unlimited.

The authors would like to thank the following people for their useful comments and suggestions regarding earlier versions of this paper: John Anderson, Pat Carpenter, Janet Davidson, Kevin Dunbar, David Klahr, Ken Koedinger, Ken Kotovsky, Deepak Kulkarni, Jay McClelland, Allen Newell, Mark St. John, David Steier, Ross Thompson, and Gregg Yost.

Correspondence and requests for reprints should be sent to Herbert A. Simon, Department of Psychology, Carnegie-Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213.